

## A NEW HYBRID TECHNIQUE FOR MODELING DENSE STAR CLUSTERS

CARL L. RODRIGUEZ<sup>1,2</sup>, BHARATH PATTABIRAMAN<sup>1,2,3</sup>, SOURAV CHATTERJEE<sup>1,2</sup>, ALOK CHOUDHARY<sup>1,3</sup>, WEI-KENG LIAO<sup>1,3</sup>,  
 MEAGAN MORSCHER<sup>1,2</sup>, FREDERIC A. RASIO<sup>1,2</sup>

*Draft version November 4, 2015*

### ABSTRACT

The “gravitational million-body problem,” to model the dynamical evolution of a self-gravitating, collisional  $N$ -body system with  $N \sim 10^6$  over many relaxation times, remains a major challenge in computational astrophysics. Unfortunately, current techniques to model such a system suffer from severe limitations. A direct  $N$ -body simulation with more than  $10^5$  particles can require months or even years to complete, while an orbit-sampling Monte Carlo approach cannot adequately treat the details of the core dynamics, particularly in the presence of many black holes. We have developed a new technique combining the precision of direct  $N$ -body codes with the speed of a Monte Carlo approach. Our Rapid And Precisely Integrated Dynamics code, the **RAPID** code, statistically models interactions between neighboring stars and stellar binaries while integrating directly the orbits of stars in the cluster core. This allows us to accurately simulate the dynamics of the black holes in a realistic globular cluster environment without the burdensome  $N^2$  scaling of a full  $N$ -body integration. We compare models of idealized globular clusters created by the **RAPID** approach to direct  $N$ -body and Monte Carlo models. Our tests show that **RAPID** can reproduce the half-mass and core radii of the direct  $N$ -body models far more accurately than the Monte Carlo approach and in  $\sim 1/200th$  of the computing time. With this technique, it will be possible to create realistic models of Milky Way globular clusters with sufficient rapidity to explore the full parameter space of dense stellar clusters.

*Keywords:* keywords

### 1. INTRODUCTION

The dynamics of dense star clusters is one of the most challenging problems of modern computational astrophysics. The large number of particles, high interaction rate, and large number of physical processes with vastly different physical timescales conspire to make globular clusters (GCs) and galactic nuclei (GN) uniquely difficult to model. In particular, the large number of black holes (BHs) in both GCs and GN often dynamically interact on a significantly shorter timescales than the rest of the cluster (Spitzer 1969). Although only comprising a small fraction of the total cluster mass, these BHs provide the dominant energy source for GCs, especially after the BH-driven core-collapse (Morscher et al. 2015). Hence, understanding their dynamics is critical to understanding the overall evolution and present-day appearance of these systems (Mackey et al. 2008). Unfortunately, since the orbital and interaction timescales of these BHs are frequently orders-of-magnitude smaller than the interaction timescale of a typical star in the cluster, resolving these effects can be particularly difficult.

Modern stellar dynamics codes have intensely investigated GCs, with the majority of work focusing on two approaches. The  $N$ -body approach directly integrates the force of every particle on every other particle, with the current generation of codes (Portegies Zwart et al. 2001; Harfst et al. 2008; Nitadori & Aarseth 2012; Capuzzo-Dolcetta et al. 2013; Wang et al. 2015) making extensive

use of state-of-the-art hardware acceleration and algorithmic enhancements. While extremely precise, this approach can require years (e.g. Heggie 2014) to complete a full simulation of a realistic Milky-Way GC.

As such, an approximate Monte Carlo (MC) technique is often used in place of a full direct summation (Hénon 1971; Giersz 1998; Joshi et al. 2000; Freitag & Benz 2001). Whereas an  $N$ -body approach computes the orbit of stars directly, the orbit-sampling approach assumes that particle orbits remain fixed on a dynamical timescale, only changing due to slight perturbations from two-body encounters between neighboring particles. This allows the orbits to be sampled statistically, and since computing a single orbit in a fixed spherical potential is significantly faster than computing the precise orbits in the full- $N$  potential of a cluster, these MC models can be generated in at most a few days or weeks. However, the assumptions of spherical symmetry and dynamical equilibrium break down in the BH-dominated core, where the potential and the particle orbits are primarily determined by a small number of particles. In practice, this means that the core radius of a MC model significantly disagrees with the core radius of the same model integrated with the more-precise direct  $N$ -body approach (Morscher et al. 2015).

GCs are formed as the result of a burst of star formation in the early universe. Approximately 10 to 20 Myr after this formation is complete, the most massive stars in the cluster collapse, yielding hundreds to thousands of BHs (Belczynski et al. 2006). As the BHs are more massive than the typical cluster star, they are rapidly driven to the center of the GC by dynamical friction (Fregeau et al. 2002); once there, the number density of BHs is sufficient to form binaries via three-body encoun-

<sup>1</sup> Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA)[e-mail: [cr@u.northwestern.edu](mailto:cr@u.northwestern.edu)]

<sup>2</sup> Dept. of Physics and Astronomy, Northwestern University, 2145 Sheridan Rd, Evanston, IL 60208, USA;

<sup>3</sup> Dept. of Electrical Engineering and Computer Science & Northwestern University, Evanston, IL, USA

ters. While it was long-assumed that these BHs would not be retained in GCs to the present day (e.g. Sigurdsson & Hernquist 1993), recent evidence has begun to suggest otherwise.

The past decade has seen the detection of the first BH in an extragalactic GC by Maccarone et al. (2007) and several recent detections in Milky Way GCs, (Chomiuk et al. 2013; Strader et al. 2013; Miller-Jones et al. 2015), including two BH candidates in M22 (Strader et al. 2012). These observational results complimented recent theoretical results suggesting that the GCs can potentially retain hundreds of BHs up to the present day (Morscher et al. 2015). The importance of BHs in GCs cannot be overstated. In addition to determining the structural and evolutionary properties of the clusters, GCs also have important implications for BH astrophysics. GCs can produce X-ray binaries at a significantly higher rate than the galactic field (Clark 1975), suggesting that there might be  $\sim 100$ s of low-mass X-ray binaries in Milky Way GCs (Pooley et al. 2003). Furthermore, recent studies have shown that the second-generation of gravitational-wave detectors can potentially detect  $\gtrsim 100$  binary BH mergers per year from binaries forged in the cores of GCs (Rodriguez et al. 2015; Antonini et al. 2015). As such, understanding the dynamics of these systems is critical.

What is needed is a technique that combines the speed of the MC approach with the precision of a direct  $N$ -body integration. In this paper, we describe a new code, the *Rapid and Precisely Integrated Dynamics* (RAPID) code, which combines both methods into a “best of both worlds” approach. In this method, the majority of particles are modeled with our parallel Hénon-style code, the Cluster MC (CMC) code (Pattabiraman et al. 2013), while the orbits of BHs are integrated directly with the Kira  $N$ -body integrator (Portegies Zwart et al. 2001). We find that this technique accurately reproduces the core radii and BH dynamics of a full direct  $N$ -body integration, with similar runtimes to the MC approach. Although we only integrate the BH orbits directly in the current work, the method is general, allowing us to select any population of particles in the cluster for  $N$ -body integration.

In Section 2, we briefly review the  $N$ -body and MC approaches, and describe the combination of the two approaches as implemented in RAPID code. In Section 3, we describe a single RAPID timestep, illustrating the technical details of the approach, while in Section 4, we describe the parallelization strategy that allows us to compute particle positions and velocities via orbit sampling and direct  $N$ -body simultaneously. Finally, in Section 5, we compare the properties of four idealized GCs as modeled by NBODY6, CMC, and RAPID. Throughout the paper, we will frequently refer to the “stars” and “BHs” in the cluster separately. In our current method, the stars are modeled with CMC and the BHs are integrated with Kira. This shorthand is to delineate which systems are being modeled by which technique, even though the particles under consideration are point-mass particles.

## 2. HYBRIDIZATION APPROACH

In this section, we provide a brief overview of the current methods employed to model GCs, and describe how our approach combines the virtues of both methods. Both the  $N$ -body and MC approaches are the result of

decades of precision work by multiple groups. For a more comprehensive description of collisional  $N$ -body dynamics, see Aarseth (2003) or Dehnen & Read (2011). A review of MC methods can be found in Freitag (2008).

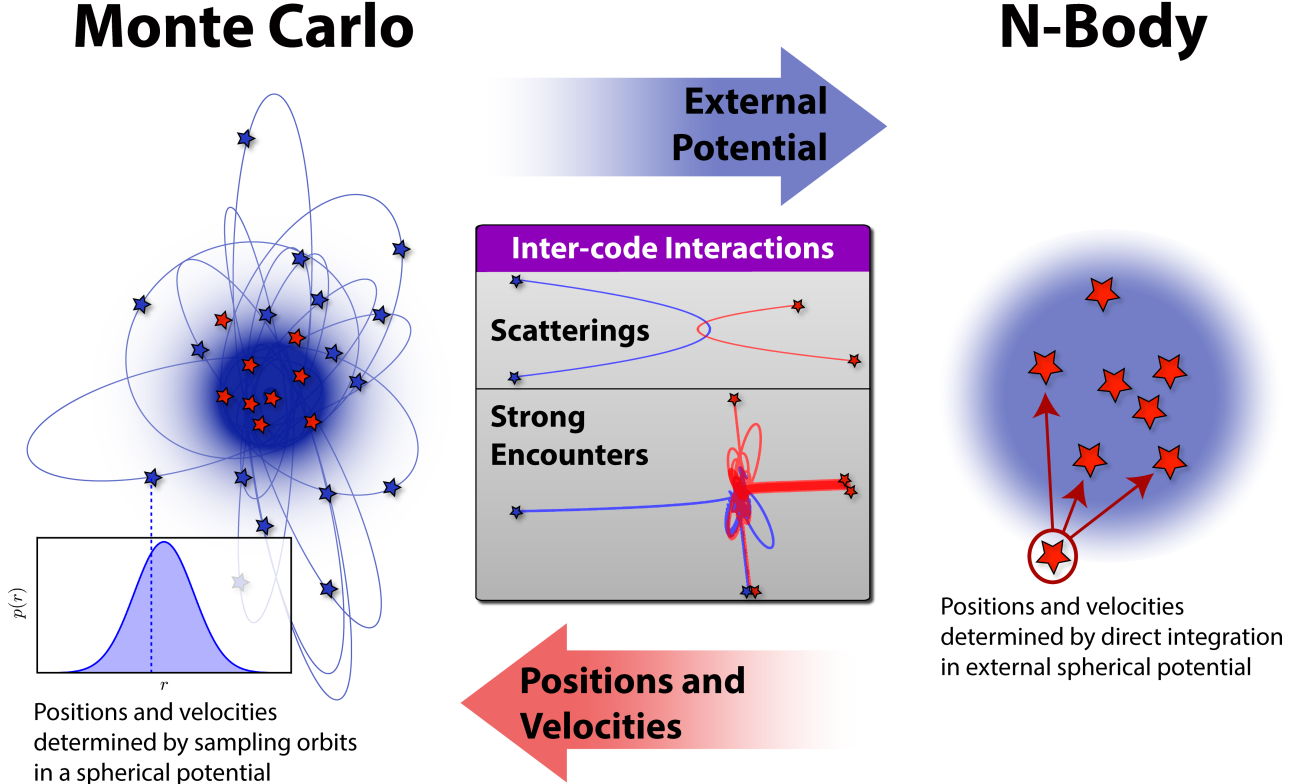
It should be noted that RAPID is not the first attempt at a hybrid  $N$ -body/statistical sampling approach to stellar dynamics. In particular, the hybrid approach developed by McMillan & Lightman (1984b) combined a Fokker-Planck sampling code with a direct  $N$ -body approach, in order to study GCs undergoing core collapse (McMillan & Lightman 1984a; McMillan 1986). The RAPID code continues this tradition of attempting to “have it all”, by combining the best of the direct integration and statistical sampling methods.

### 2.1. Direct $N$ -Body Integration

The physical principle behind a direct  $N$ -body integrator is simple: since the force on any given particle is the sum of the gravitational force from every other particles in a given system, the most accurate way to model such a system is to numerically sum all the forces. This is the underlying principle behind the  $N$ -body integrators. The most frequently used of these codes, the NBODY series of codes, have been improved and finely tuned with additional physics, including stellar evolution (Hurley et al. 2001), algorithmic regularization (Aarseth 1999), post-Newtonian chain regularization, (Aarseth 2012), and GPU acceleration (Nitadori & Aarseth 2012). With advanced hardware and a minimal number of simplifying assumptions, direct integration is the most precise method available for modeling dense stellar systems.

However, this precision comes at a cost. Since a direct summation scales as  $N^2$  in the number of particles, the time requirements for a simulation with  $N \gtrsim 10^6$  are exceedingly restrictive. The largest simulation attempted with NBODY6 is currently the  $N = 5 \times 10^5$  model of galactic GC M4 performed by Heggie (2014), requiring 2.5 years on a dedicated GPU system. More recently, the current state-of-the-art parallelized code NBODY6++GPU (Wang et al. 2015) can model a realistic ( $N = 10^6$ ) cluster in little more than a year. Despite these remarkable achievements, simulation times in excess of  $\sim 1$  year for large systems preclude any reasonable exploration of the parameter space of initial conditions of GCs, and any collisional models of GN ( $N = 10^7 - 10^9$ ) remain beyond the capabilities of the current generation of direct summation techniques. To answer astrophysical questions related to such systems, a more rapid technique is called for.

For our hybrid approach, we use the Kira  $N$ -body integrator, included as part of the Starlab software package (Portegies Zwart et al. 2001). Like the NBODY series of codes, Kira is a 4<sup>th</sup>-order Hermite predictor-corrector integrator with a block timestep scheme. Kira also integrates close encounters and tightly-bound multiples using Keplerian regularization, where sufficiently-isolated hyperbolic and tightly bound binaries are evolved as analytic two-body systems. Additionally, Kira organizes its internal data using easily-modifiable C++ class structures, and includes an easily-customizable module for including an external gravitational potential. These two features make it ideal for inclusion in the hybrid method.



**Figure 1.** Cartoon explanation of the RAPID Code. The majority of particles are evolved via the MC approach, with their positions and velocities determined by randomly sampling orbits in the spherical potential of the entire cluster. Meanwhile, a small handful of the most massive objects (the BHs) have their positions and velocities determined by direct  $N$ -body integration inside the external potential of the MC stars. Particles from both codes can interact via two-body encounters and strong scattering encounters, allowing energy to flow between the two systems.

## 2.2. Orbit-Sampled Monte Carlo

The MC approach assumes that the large-scale evolution of a star cluster can be modeled as slow transitions from one equilibrium configuration to the next. These transitions are driven by two-body scatterings (relaxations) between particles in the cluster. Because it is the statistically-predictable effect of many of these two-body scatterings that let energy flow through the cluster, one can describe the cumulative effect of these relaxations as a single, effective two-body encounter. We often describe these systems in terms of their dynamical timescales (the crossing time for a single particle) and their relaxation timescales (the average time for the velocity of a single particle to change significantly). For a system in dynamical equilibrium:

$$T_{\text{relax}} \approx \frac{0.1N}{\log N} T_{\text{dyn}} \gg T_{\text{dyn}} \quad (1)$$

where  $N$  is the number of particles (see Binney & Tremaine 2011). With a sufficiently large  $N$ , a star cluster can be thought of as a series of independent orbits that only change on a  $T_{\text{relax}}$  timescale. This assumption eliminates the need to directly compute the forces upon a single particle on an orbital timescale. Instead, we only need to determine the shape of a star's orbit after it has changed velocity due to encounters with other stars on the relaxation timescale.

For some cases, such as a spherically asymmetric mass

distribution, the orbit of the particle must be integrated numerically (e.g. Vasiliev 2014; Vasiliev et al. 2015); however, for most applications to large collisional star clusters (such as GCs and GN) the background gravitational potential can be assumed to be spherical. This allows the clever theorist to determine a star's position and velocity by analytically sampling a random point along its orbit. This *orbit-sampling MC* approach, first developed by Hénon (1971) and built upon by multiple groups (Stodoikiewicz 1982; Giersz 1998; Joshi et al. 2000; Freitag & Benz 2001) can model stellar systems with  $N \gtrsim 10^7$  particles in a fraction of the time of a direct  $N$ -body simulation. Since the primary computational bottleneck is the sorting of particles by radius, the MC method scales as  $N \log N$  in the number of particles, which is a significant improvement over the  $N^2$  scaling of a direct  $N$ -body integration. Combined with the use of a relaxation timestep (which gives an additional  $N/\log N$  speed increase over the dynamical timestep of the  $N$ -body approach), the MC method can easily model large systems that are simply beyond the reach of other techniques.

However, the assumptions that enable the speed of the MC method can easily break down in some of the most interesting regions of parameter space. Once mass segregation is complete, the evolution of a GC is largely determined by the small number of BHs that have accumulated in the core. This can consist of as few as  $\sim 100$  BHs, and is often accompanied by further deep collapses

involving only  $\sim 5$  BHs (Morscher et al. 2015). Since the dynamics of these small, spherically asymmetric systems change rapidly on an orbital timescale, the MC method is unable to accurately follow the evolution of these BHs in the center of the cluster. And since this small cluster of BHs forms the hard binaries whose binding energy acts as a power source for the entire cluster, their dynamics *must* be accurately modeled to understand the long-term evolution of the cluster.

Our orbit-sampling Cluster MC code, **CMC**, was first developed by Joshi et al. (2000), based on the original developments by Hénon (1971) and Stodoikiewicz (1982). As the code considers interactions between individual stars, **CMC** incorporates multiple physical processes, including stellar evolution (Hurley et al. 2000, 2002), strong three-body and four-body scatterings with the small- $N$  integrator **Fewbody** (Fregeau et al. 2004), probabilistic three-body binary formation (Morscher et al. 2013), and physical collisions. Additionally, **CMC** has recently been parallelized to run on an arbitrary number of computer processors (Pattabiraman et al. 2013). This MPI parallelization makes **CMC** an ideal code base for **RAPID**, as the current parallelization scheme can be easily expanded to allow the  $N$ -body integration to run in parallel to the MC

### 2.3. Hybrid Partitioning

The cornerstone of any hybrid modeling technique is the domain decomposition between methods. Since it is the BHs that are driving the non-spherical, non-equilibrium dynamics in the cluster core, the natural division is for Kira to integrate the BHs while **CMC** integrates all the remaining stars in the cluster. By default, **RAPID** divides the system using one of two criteria:

- a **mass criterion**, which divides the system according to a specified threshold, where particles above the threshold are considered BHs and particles below it are considered stars, and
- a **stellar evolution criterion**, in which objects identified as BHs by stellar evolution are integrated by Kira, and all other objects are integrated by **CMC**

By default, **RAPID** employs the first criterion for point-mass simulations (with a user-specified threshold mass), and the second criterion for simulations using stellar evolution. Any mixed objects (e.g. a BH-star binary) are evolved in **CMC**, in order to treat the binary stellar evolution consistently.

The **RAPID** code builds upon the **CMC** parallelization described in Pattabiraman et al. (2013), and is designed to be run on a distributed computational system with at least 2 parallel MPI processes: one for the MC integration, and one for the  $N$ -body integration. An initial **RAPID** run begins with all objects integrated with **CMC** on all available processes. After a user-specified criterion is met, a single MPI process initializes the Kira  $N$ -body integrator. At this point, any stars that are on that **CMC** process are transferred to the remaining **CMC** process(es), and the BHs are collected on the Kira process for the  $N$ -body integration. Once both integrators have been initialized, particles can be transferred back and forth between the Kira and **CMC** processes (e.g. a star becomes a BH through stellar evolution, and is copied from

**CMC** to Kira). Additionally, information about the particles must frequently be communicated back and forth between **CMC** and Kira. At each timestep, **CMC** needs to know the positions and velocities of the BHs, while Kira needs to know the external potential of the stars. See Section 4 for the details of the parallelization strategy. At the end of each timestep, the information from Kira and **CMC** is combined and printed to file, in order to create a coherent cluster model. See Figure 1.

## 3. RAPID TIMESTEP

At its core, the main difference between **RAPID** and **CMC** is the computation of the orbits and positions of the BHs. Here, we describe in detail a single **RAPID** timestep, highlighting the differences between the hybrid approach and a standard **CMC** integration. See Figure 2.

### 3.1. Compute the Potential

Since the cluster is assumed to be spherically symmetric, the gravitational potential at radius  $r$  between stars  $k$  and  $k+1$  can be expressed as:

$$\Phi(r) = G \left( -\frac{M_k}{r} - \sum_{i=k+1}^N \frac{m_i}{r_i} \right) \quad (2)$$

where  $M_j \equiv \sum_{i=1}^j m_i$ . As the stars are radially sorted, the potential at each star can be computed recursively with a single pass from the outermost star inwards:

$$\Phi_{N+1} = 0 \quad (3)$$

$$\Phi_k = \Phi_{k+1} - GM_k \left( \frac{1}{r_k} - \frac{1}{r_{k+1}} \right) \quad (4)$$

In the **RAPID** code, two potentials are calculated: the full spherical potential,  $\Phi$ , of all stars and BHs, and a MC-only potential,  $\Phi^{\text{MC}}$ , computed only with the stars in **CMC**. The MC potential is sent from the **CMC** processes to the  $N$ -body process, and used as an external potential for the Kira integration.

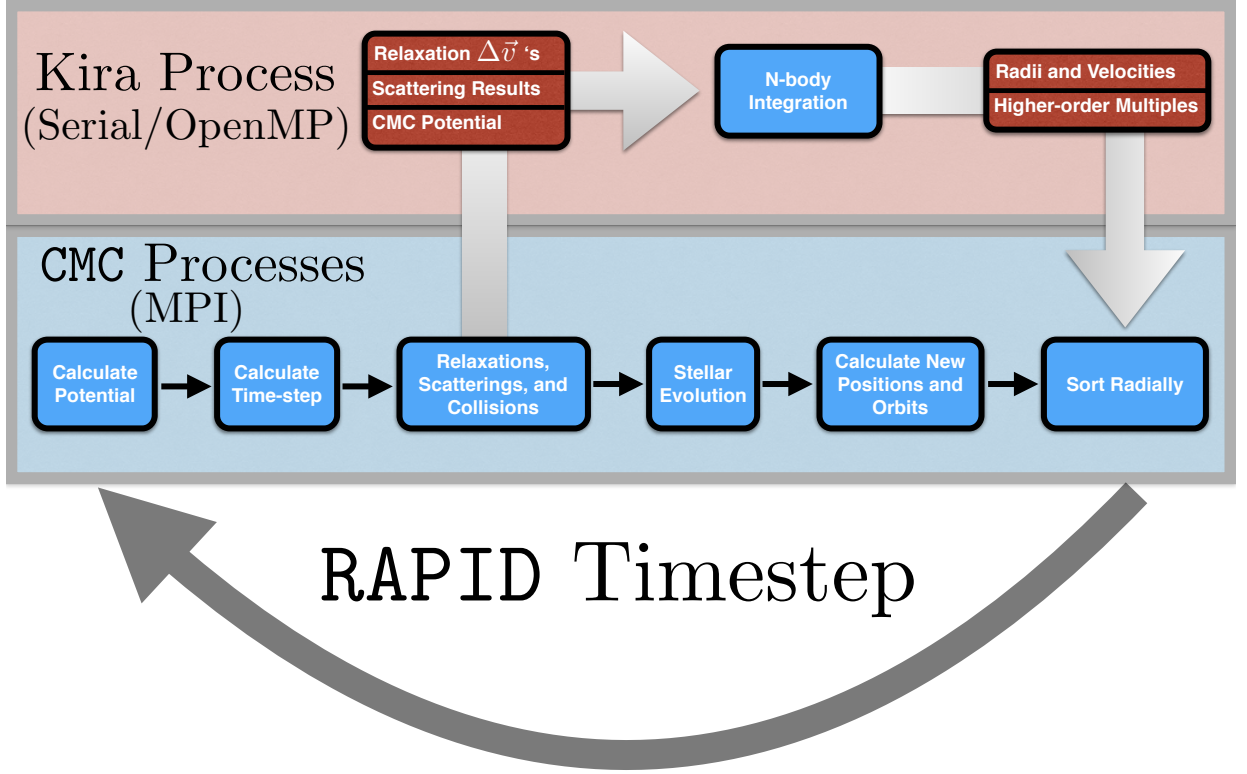
### 3.2. Select the Timestep

Since **CMC** and **RAPID** consider multiple physical processes, the timestep selection must consider multiple relevant timescales: the timescale for two-body relaxation ( $T_{\text{rel}}$ ), the timescales for binary-single and binary-binary strong scatterings ( $T_{\text{BS}}$  and  $T_{\text{BB}}$ ), the timescale for physical collisions ( $T_{\text{coll}}$ ), the timescale for stellar evolution to contribute a significant change in an objects mass ( $T_{\text{SE}}$ ), and the timescale for tidal stripping of stars to significantly alter the cluster mass ( $T_{\text{tid}}$ ). The simulation timestep is then selected to be the minimum of each of these timescales, or

$$\Delta T_{\text{CMC}} = \min(T_{\text{rel}}, T_{\text{BS}}, T_{\text{BB}}, T_{\text{coll}}, T_{\text{SE}}, T_{\text{tid}}) \quad (5)$$

The **RAPID** timestep is chosen in a similar fashion. The timescale for each physical process is computed for all particles in the cluster (stars and BHs), and the minimum (or a fraction of the minimum) is selected as the current timestep. In Kira, this timestep determines how many dynamical times the  $N$ -body system will be advanced.





**Figure 2.** Flowchart of the RAPID Code. After the hybridization is initialized, the BHs are transferred to the Kira process, while the remaining Monte Carlo processes sample the orbital positions of the remaining stars. The communication between the Kira process and the CMC process(es) is handled by an extension of the MPI scheme developed in CMC (Pattabiraman et al. 2013), while the *N*-body process can be run in serial or with OpenMP threading.

Unlike CMC, RAPID can produce BH multiples with an arbitrary number of components. To compute the interaction timescale for scatterings between stars and BH multiple systems, the timestep is chosen using the same prescription as the  $T_{BS}$  and  $T_{BB}$  timesteps, but with the semi-major axis of the outermost binary pair as the effective width of the system.

### 3.3. Perform Dynamical Interactions

After selecting the global simulation timestep, the CMC processes apply the standard nearest-neighbor interactions (two-body relaxations, strong encounters, and collisions) to successive pairs of particles in the radially sorted array of all cluster particles. The interactions are identical to those employed in our previous CMC studies. What differs in the RAPID approach is how the outcomes of interactions between CMC stars and Kira BHs are handled. Specifically:

- **Star-Star** interactions are handled in the same fashion as in a pure-CMC integration (see Pattabiraman et al. 2013, and references therein).
- **Star-BH** interactions are also handled in the same fashion (by CMC); however, in addition to updating the local BH information stored in CMC, the dynamical changes to each particle are communicated back to the Kira process once the interaction step is complete.
- **BH-BH** nearest-neighbor interactions are skipped,

since such encounters will be performed with greater accuracy in the *N*-body integration.

In CMC, two-body relaxations are performed by setting nearest-neighbor stars along randomly selected hyperbolic orbits that are consistent with their radial and tangential velocities. The hyperbolic encounter modifies the velocities of both stars, allowing for an exchange of energy and angular momentum. At the end of the timestep, we communicate the full 3D velocity changes for each BH to the Kira process. The velocity of each particle in the *N*-body is updated by adding the full-3D velocity perturbation to previous 3D velocity of the BH.

In addition to two-body relaxations, CMC integrates strong scattering encounters between neighboring multiples (such as binary-single neighbors or binary-binary neighbors) using the *Fewbody* small-*N* integrator (Fregeau et al. 2004). In RAPID, we also allow for strong encounters between neighboring stars and BHs. However, as Kira can produce BH higher-order multiple systems, (triples, quadruples, etc), we have modified *Fewbody* to perform scatterings between systems of arbitrary size, such as single-multiple and binary-multiple encounters. We ignore multiple-multiple scatterings, since CMC does not track higher-order stellar triples, and BH-BH encounters are performed naturally in Kira.

Once *Fewbody* has completed the scattering (which can take several seconds of CPU time for compact higher-order multiple systems) the output is then sent back to CMC and Kira separately. For higher-order multiples, the full 3D position and velocity of each BH component, rel-

ative to the multiple center-of-mass, is sent to the Kira process. The multiple is then reinserted into the  $N$ -body integration at its previous position.

The only exception to this procedure is the formation of mixed-multiple systems (a binary or higher-order multiple with both BH and stellar components). We evolve any BH-Star binaries in CMC. For higher-order multiples with both stellar and BH components, the multiple is hierarchically broken into binary and single stars until all stars are found in single and binary objects (to be evolved by CMC), while the remaining BH objects are sent to Kira. The binding energy from any broken multiples is subtracted from the center-of-mass kinetic energy of the (now unbound) components.

### 3.4. Perform Stellar Evolution

RAPID considers realistic stellar evolution using the Single Stellar Evolution (SSE) and Binary Stellar Evolution (BSE) packages of Hurley et al. (2000, 2002). This is identical to the previous implementation in CMC. No stellar evolution is required by the direct  $N$ -body, since the only particles integrated by Kira are BHs. As stated above, all mixed BH-Star objects are integrated in CMC. This is to ensure that the binary stellar evolution for BH-star systems is performed consistently.

### 3.5. Calculate New Orbits and Positions

After the dynamical information for each particle has been updated, a new orbital position and velocity, consistent with the particle's new energy and angular momentum, must be computed. Since the dynamical state of each particle is up-to-date in CMC and Kira, the MC and  $N$ -body integrations can be performed in parallel by their respective processes.

#### 3.5.1. Orbit Calculation (MC)

Orbits in CMC are determined using the standard Hénon-style orbit-sampling approach. We assume that the orbit is entirely a function of the star's kinetic energy and the spherical potential of the cluster. We first constrain the radial extent of the orbit by computing the two zeros ( $r_{\min}$  and  $r_{\max}$ ) of the energy equation:

$$2E - 2\Phi(r) + J^2/r^2 = 0 \quad (6)$$

Then, we probabilistically select a new radius for the star based on the star's new orbit. Since the probability of finding the star at a given radius is proportional to the time the particle spends at that radius, we can express the probability of finding the star at a specific radius as

$$P(r)dr = \frac{dt}{T} = \frac{dr/|v_r|}{\int_{r_{\min}}^{r_{\max}} dr/|v_r|} \quad (7)$$

We then draw a random sample from  $P(r)$ , and use the value as the particle's position for the next timestep.

#### 3.5.2. Orbit Calculation ( $N$ -body)

Because we have dynamically perturbed the BHs though scattering and a new external potential, the gravitational force and its higher derivatives must be recalculated before resuming the Kira integration. Otherwise, the dynamical changes to the BH velocities would produce discontinuities in the higher derivatives of the force,

breaking the smoothness needed for the 4<sup>th</sup>-order Hermite integrator to work. The  $N$ -body system is reset using Kira's built-in reinitialization function, which recomputes the acceleration and jerk for each BH explicitly. The position and velocities of all the particles (up to any changes from dynamical interactions) are not modified.

Once the system has been reinitialized, the  $N$ -body configuration is directly integrated with the Kira integrator for a number of dynamical times equal to the current RAPID timestep. For each particle, the total force is computed as the sum of the external force, from  $\Phi^{\text{MC}}$ , and the internal force computed via direct summation of the other BHs. We apply the external CMC potential to each isolated BH and the center-of-mass of each bound multiple system by computing the acceleration and jerk from the external potential

$$\begin{aligned} \vec{a}_{\text{ext}} &= -\frac{\partial \Phi^{\text{MC}}}{\partial r} \hat{r} \\ \vec{j}_{\text{ext}} &= -\frac{1}{r} \frac{\partial \Phi^{\text{MC}}}{\partial r} (\vec{v} - (\vec{v} \cdot \hat{r})\hat{r}) - \frac{\partial^2 \Phi^{\text{MC}}}{\partial r^2} (\vec{r} \cdot \hat{r})\hat{r} \end{aligned}$$

and adding this to the acceleration and jerk of each BH using Kira's external potential module. To simplify the calculation of the potential and its derivatives, we select 30 stars evenly in  $\log r$  from the innermost CMC star to the stripping radius of the  $N$ -body simulation, and copy their radii and  $\Phi^{\text{MC}}$  values to the  $N$ -body integrator. The values of  $\Phi^{\text{MC}}$  and its radial derivatives are then calculated via 5<sup>th</sup>-order Lagrange Polynomial interpolation. This technique is inspired by the radial potential sampling used in Vasiliev (2014).

### 3.6. Sort Radially

Finally, once all the relevant physics has been applied and the data collected on the CMC processes, the particles must be sorted in order of increasing radial distance from the GC center. The sorting is performed in parallel by all CMC processes using the parallel Sample Sort algorithm described in Pattabiraman et al. (2013).

## 4. PARALLELIZATION STRATEGY

To incorporate the Kira integrator into CMC, we make the following modifications to our parallelization strategy. When the simulation starts, all particles are evolved using the CMC scheme. As described in Section 2.3, the activation criterion for the  $N$ -body integrator depends on the type of particles being integrated. For point particle simulations, the  $N$ -body integrator is begun immediately, whereas for star clusters modeled with stellar evolution, the  $N$ -body integrator is only activated once a certain number ( $\sim 25$ ) of BHs have been formed. Once the activation condition is met, we divide the entire set of particles into two sets: MC stars and  $N$ -body BHs.

At the same time, we divide the existing set of  $p$  processes into two separate groups: a single Kira process for integrating the BHs, and  $p-1$  CMC processes for integrating the stars. Any MPI communication is handled by two custom intracommunicators: one corresponding to all  $p$  processes and one restricting communication to the  $p-1$  CMC processes. The latter allows us to employ the same parallelization strategy described in Pattabiraman et al. (2013) with minimal modification. When the processes

Model Name	Initial Conditions			Runtimes (Hours)		
	$M_{\text{BH}}/M_{\text{star}}$	$m_{\text{BH}}/m_{\text{star}}$	$T_{\text{end}} (t_{\text{dyn}})$	CMC	NBODY6	RAPID
64k-0.01-10	0.01	10	14000	2	425	2
64k-0.02-10	0.02	10	20000	3	304	2
64k-0.01-20	0.01	20	10000	2	554	1
64k-0.02-20	0.02	20	20000	2	512	2

**Table 1**

The model names, initial conditions, and runtimes for each of the four idealized GC models. The runtimes quoted are the relevant walltimes for each method. The **CMC** and **RAPID** models were run on 4 Intel Xeon E5-2670 Sandy Bridge processors, while the **NBODY6** models were run on 8 processors and 1 Nvidia Tesla M2090 GPU.

are split, all the BHs are sent to the Kira process, where their coordinates are converted from  $(r, v_r, v_t)$  space to the full 6D phase space by randomly sampling the orientation of the position and velocity vectors. This sets the initial conditions for the  $N$ -body integration. In addition, the Kira process maintains a radially-sorted array of  $(r, v_r, v_t)$  for all BHs. This is done to facilitate easier communication with **CMC** and to minimize the required MPI communication between **CMC** and Kira every timestep. Although the **CMC** processes hand over their BHs to the Kira process, the BHs are not deleted from their local arrays. The BHs are left intact yet inert, so that their positions and velocities can be updated upon completion of the  $N$ -body integration.

Once the **CMC** and Kira MPI processes have been initialized, the hybrid method must allow both codes to interact while minimizing the amount of MPI communication. This is accomplished by a series of intermediate arrays, designed to store and transmit the minimum amount of information back and forth between the **CMC** and Kira. Communication only occurs at two points during a **RAPID** timestep. The first occurs after the relaxation and strong-encounters have been performed, in order to communicate the dynamical changes from **CMC** to Kira. The second occurs after both systems have completed their respective orbit computations, to communicate the new dynamical positions and velocities from Kira back to **CMC**.

#### 4.1. *CMC to Kira Communication*

After the **CMC** processes have computed a new potential, performed two-body relaxations, and any strong encounters, the dynamical changes to the BHs must be communicated to the Kira process. This is done with three distinct communications:

- The **MC potential**,  $\Phi^{\text{MC}}$ , is sent to the Kira process as two arrays, containing the radius and cluster potential of every star (excluding the BHs) in **CMC**. The Kira process selects 30 of these stars as described in Section 3.5.2, and passes the information to the Kira integrator to use when computing the external force.
- The **two-body relaxations** are communicated as an array of objects, each containing a particle ID and a 3-dimensional  $\Delta\vec{v}$ . These weak velocity perturbations are added to the single BHs and the centers-of-mass of any BH multiple systems before the  $N$ -body system is reinitialized by Kira.
- The results of **strong encounters** are communicated differently depending on the type of encounter. For binary BHs that have experienced a

strong encounter with a star, only the change in semi-major axis and eccentricity are communicated back to Kira. For triples and higher-order multiples, the full position and velocity of every BH in the multiple is communicated to Kira. To reduce communication, any hierarchical information is not transmitted, and the Kira process reconstructs the hierarchy locally before the  $N$ -body system is reinitialized.

Since all the information previously described does not significantly modify the radial positions or velocities of the particles in the  $N$ -body (by assumption, the MC approach requires that  $\Delta v/v \ll 1$ ), the dynamical state of the  $N$ -body system is preserved between **RAPID** timesteps. The one exception is strong encounters in which a single bound BH multiple is broken into components. Since strong encounters in **CMC** are performed by assuming the scatterings are isolated from the cluster potential at infinity, the resultant components cannot be placed at the same infinite location in the  $N$ -body system. For such systems, the components are placed at the correct radius and random orientations, similar to the initial transfer of BHs from **CMC**.


Finally, we allow for the possibility that **CMC** may create new BHs, either through stellar evolution or through strong encounters which produce single or binary BHs. We add any such new BHs to Kira. The new BH is then flagged as a BH in the local array of the **CMC** process that created it, to ensure it is not evolved by **CMC** during the next timestep.

#### 4.2. *Kira to CMC Communication*

After the Kira integrator has computed the orbits of the BHs, the new dynamical state must be communicated back to the **CMC** processes. First, the dynamical 6-D phase space information for each particle in the  $N$ -body is projected back to the reduced  $(r, v_r, v_t)$  basis, and copied back into place in the intermediate star array on the Kira process. The intermediate BH array is then divided up and communicated back to the respective **CMC** processes.

In addition to the positions and velocities, any new objects, such as single BHs, newly formed binaries, or higher-order multiples, are communicated as new objects to the last  $(p-1)$  **CMC** processes, to be placed in the correct process once the **CMC** (and intermediate BH) arrays are sorted. For single and binary BHs, this is accomplished by sending the usual  $(r, v_r, v_t)$  for each system and the semi-major axis and eccentricity for any binaries to the **CMC** processes. For higher-order multiples, the full 6D dynamical information is transmitted back using the

**Figure 3.** The half-mass and core radii for all four models as determined by direct  $N$ -body (NBODY6, in red), the Monte Carlo (CMC, in blue), and the hybrid approach (RAPID, in purple). Although the half-mass radii agree reasonably well (within statistical fluctuations) between all three codes, the RAPID code replicates the core radii of the direct  $N$ -body method significantly better than CMC, particularly near core collapse. At late times, the core radii begin to diverge; this begins to occur when the majority of the BH system has been ejected, and the CMC relaxation dynamics dominate the BHs. RAPID reverts to pure CMC when the number of BHs drops below a certain threshold ( $\sim 5$  BHs, indicated by the black dashed lines).



same array that sent the strong encounters from CMC in the first communication. Again, only the positions and velocities of the BH multiple components are communicated, so the hierarchical information is reconstructed on the local CMC process after communication.

Because the  $N$ -body integration is being performed in a lower-density environment than the full cluster, it is possible for Kira to produce pathologically wide binaries and higher-order multiples. This effect is particularly problematic at late times, where a handful ( $\sim 10$ ) of BHs can easily produce binaries with separations greater than the local inter-particle separation of stars. Since the unphysically large interaction cross-sections of these systems drastically shrink the CMC timestep, we break apart any multiple systems whose apocenter distance is greater than 10% of the local inter-particle separation of stars at that radius. The kinetic energy of the CMC stars is adjusted to ensure energy conservation. This criterion for breaking wide binaries is the same criterion used to break wide binaries produced by *Fewbody* in CMC.

## 5. COMPARISON RESULTS

To test the effectiveness of this hybrid approach, we compare our code to idealized models of GCs using similar initial conditions to Breen & Heggie (2013). We considered four clusters with 65,536 point-mass particles. The majority of particles are low-mass stars, while a small fraction of particles are high-mass BHs. We considered different mass ratios between BHs and stars

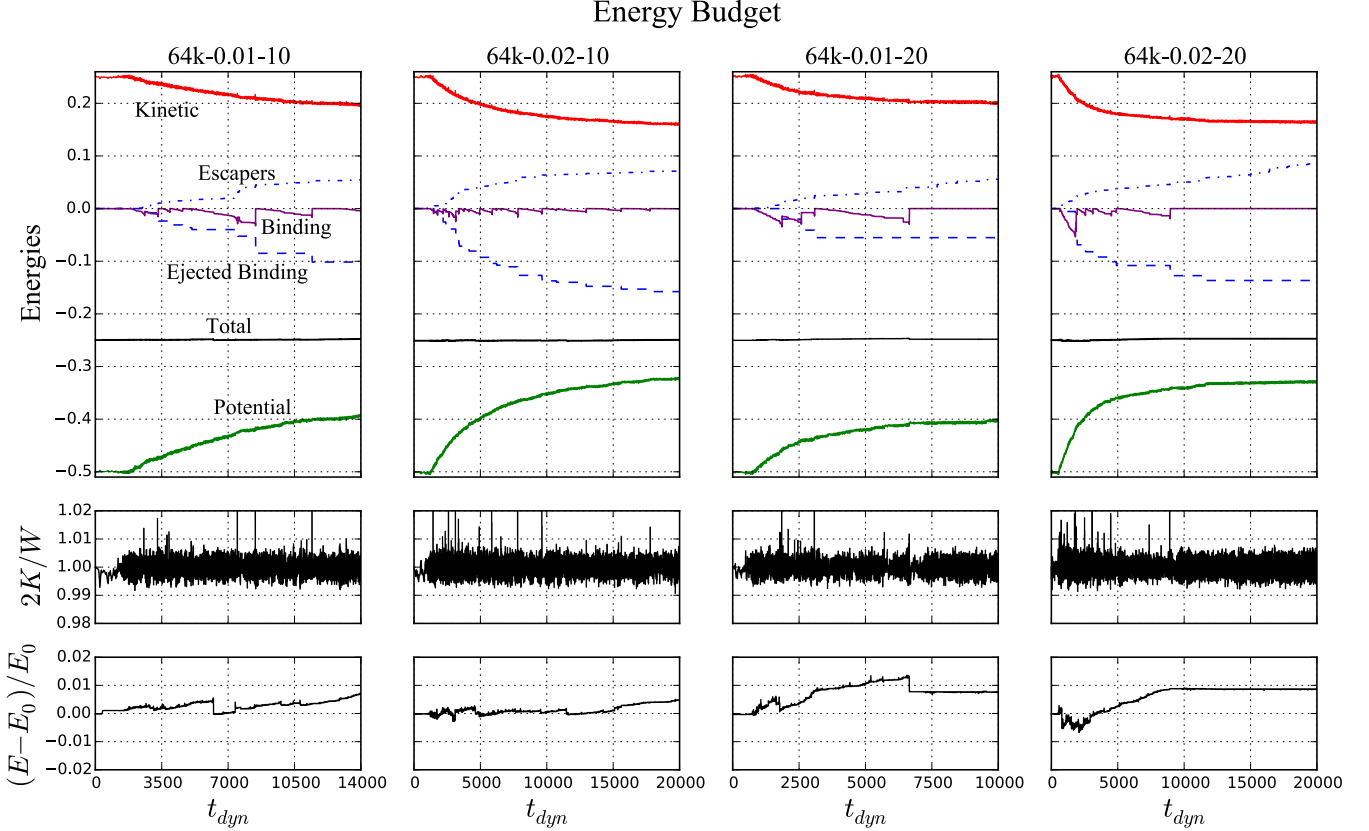
( $m_{\text{BH}}/m_{\text{star}} = 10$  and  $m_{\text{BH}}/m_{\text{star}} = 20$ ). We also varied the total mass in stars and BHs ( $M_{\text{BH}}/M_{\text{star}} = 0.01$  and  $M_{\text{BH}}/M_{\text{star}} = 0.02$ ). For each set of initial conditions, we compare the models produced by RAPID to models produced by CMC and by a pure  $N$ -body integration using NBODY6 (Aarseth 1999). Each model was integrated until the majority of the BHs had been ejected from the cluster. A summary of the initial conditions and the average runtimes for each approach is listed in Table 1.

### 5.1. Core and half-mass radii

For the purposes of this analysis, we assume that the models produced by NBODY6 are the “true” clusters, since a direct integration technique requires the fewest simplifying assumptions. We wish to know how well the models from our new hybrid technique match the models produced by the  $N$ -body approach. In Figure 3, we compare the half-mass radius and core radius for each model as determined by CMC, NBODY6, and RAPID. While the half-mass radii are relatively consistent across all models, the core radii are drastically different between the three methods. CMC consistently underestimates the core radius of each cluster immediately after core collapse, with the trend persisting until all BHs have been ejected from the cluster. This consistent underestimation of the cluster core radius has been observed in other studies using CMC (Morscher et al. 2015) and other orbit-sampling Monte Carlo codes. However, the core radii as determined by RAPID agree very well with the radii



**Figure 4.** The energy budget and overall energy conservation for the four RAPID models shown in Fig. 3. The top plot shows the total sum of the kinetic energy (red), potential energy (green), and multiple binding energy (purple) for all particles in the cluster over time. The energy carried away by the kinetic energy (blue-dotted) or binding energy (blue-dashed) of escaping particles is also shown. The sum of all energies is plotted in Hénon units (black) in the top plot, and as a fractional change in the bottom plot. The total energy conservation error is dominated by the evolution of the  $N$ -body particles in a strong external potential (the slow upward trend) and occasional strong-encounters and tightly-bound multiple systems (discontinuous jumps).



reported by the direct  $N$ -body. This suggests that the RAPID approach can correctly model the dynamics of the massive BHs which dominate the long-term evolution of GC cores.

### 5.2. Energy Conservation

To check the consistency of the RAPID approach, we examine the energy budget and conservation of each of the runs. We quantify the various forms of energy, including the kinetic and potential energy of all particles, the binding energy of multiples, and the total energy carried out of the cluster by ejected particles. The energies are plotted in Figure 4. In addition, we also consider the virial ratio ( $2KE/PE$ ) and the total energy over time for each system, as a diagnostic of the effectiveness of the method.

The total energy conservation of RAPID can vary significantly over a single run, and usually lies within 1% of the initial energy for the duration of the run. There are two significant sources of error which contribute to this energy flux. The first is the ability of Kira to integrate higher-order multiples and very close encounters accurately. This issue is not limited to the Kira integrator, and is one of the well-known issues common to all collisional dynamics simulations (the so called “terrible triples”). Furthermore, as the CMC potential is only applied to the center-of-mass of any multiple systems,

any tidal effects upon the multiples from the background cluster potential are not incorporated correctly. These integration errors manifest as discontinuous jumps in the overall energy conservation, which can be seen in the bottom panels of Fig. 4.

The second source of error arises from the integration of orbits in a fixed external potential. The 4<sup>th</sup>-order Hermite integration scheme, while typical for collisional stellar dynamics, is known to produce systematic energy errors when integrating many orbits in a fixed potential (Binney & Tremaine 2011; Aarseth 2003). This produces a small but systematically positive energy drift while integrating the BHs over many orbits. Furthermore, since the energy error increases with the number of timesteps taken by a particle, this drift cannot be corrected by increasing the  $N$ -body accuracy. Methods to improve this energy drift, such as modifying the Aarseth timestep criterion (Aarseth 2003) or increasing the order of the Hermite integrator (Nitadori & Makino 2008) are currently being investigated.

## 6. CONCLUSION

In this paper, we described the motivation and development of a new hybrid technique for dynamically modeling dense star clusters. By combining our Cluster Monte Carlo (CMC) code with the Kira direct  $N$ -body integrator, we are able to combine the speed of the MC ap-

proach with the accuracy of a direct summation. This hybrid code, the *Rapid and Precisely Integrated Dynamics* (RAPID) Code is designed to accurately model the non-equilibrium BH dynamics that powers the overall evolution of GCs and GNs. Given recent observational detection of BH candidates in GCs, and the importance of theoretical modeling of GC BHs to X-ray binary astrophysics (Pooley et al. 2003) and gravitational-wave astrophysics (Rodríguez et al. 2015; Antonini et al. 2015), understanding the dynamics of BHs in clusters is crucial to understanding BH astrophysics.

We found that the hybrid approach is able to replicate both the half-mass radius and the core radius for several  $N$ -body models of idealized GCs. Unlike a purely MC approach, RAPID can model the highly non-spherical and rapidly changing dynamics of the few BHs in the center of the cluster. This suggests that the RAPID approach can follow the dynamics of BH systems with comparable accuracy to a direct  $N$ -body integration, but with roughly the same integration time (with in a factor of 2) of an orbit-sampling Monte Carlo approach.

With this technique, it will be possible to explore regions of the GC parameter space that have remained outside the computational feasibility of direct  $N$ -body computations. In particular, by treating the central BH subcluster correctly, RAPID can explore regions of the GC and GN parameter space, including clusters with massive central BHs, that have previously been unexplored by direct collisional methods.

Two issues remain to be addressed. First, the Kira  $N$ -body integrator does not completely conserve energy in the presence of a large external potential. This effect is a well-known drawback of  $4^{th}$ -order Hermite integrators, will need to be addressed. Efforts are currently underway to increase the computational order of the Hermite predictor-corrector, improving both the accuracy and speed of the integration (e.g. Nitadori & Makino 2008).

Secondly, the regularization of binaries and higher-order multiples in Kira is based on Keplerian regularization for sufficiently unperturbed systems (see Portegies Zwart et al. 2001). However, this regularization does not include any tidal effects from the external Monte Carlo potential, which will significantly effect the long-term evolution of binaries retained by the cluster. Incorporation of these physical effects into the regularization scheme is currently underway.

Future work will investigate hardware acceleration of the  $N$ -body integrator. The Kira integrator is designed to run on the specialized GRAPE series of hardware, which yields significant improvements in computational speed. When combined with the Sapporo GPU/GRAPE library (Gaburov et al. 2009), Kira can be run on modern, distributed GPU systems with comparable performance to the NBODY series of codes (Anders et al. 2012). Hardware acceleration was not implemented the current RAPID version, since we have not considered systems with sufficiently large numbers of BHs for efficient GPU usage; however, the development of the Sapporo2 library (Bédorf et al. 2015) provides efficient GPU saturation for small- $N$  systems. We will explore the advantages of a RAPID integration with Sapporo2 in a future paper.

**Acknowledgments** We thank Eugene Vasiliev, Simon Portegies Zwart, Vicky Kalogera, Claude-André Faucher-Giguère, Douglas Hoggie, Philip Breen, and Fabio Antonini for useful discussions. CR was supported by an NSF GRFP Fellowship, award DGE-0824162. This work was supported by NSF Grant AST-1312945 and NASA Grant NNX14AP92G.

## REFERENCES

- Aarseth, S. J. 1999, *PASP*, 111, 1333  
 —. 2003, *Gravitational N-Body Simulations* (Cambridge: Cambridge Univ. Press)  
 —. 2012, *MNRAS*, 422, 841  
 Anders, P., Baumgardt, H., Gaburov, E., Zwart, S. P., & Portegies Zwart, S. 2012, *MNRAS*, 421, 3557  
 Antonini, F., Chatterjee, S., Rodriguez, C. L., et al. 2015, *ArXiv eprints arXiv:1509.05080*  
 Bédorf, J., Gaburov, E., & Portegies Zwart, S. 2015, *ComAC*, 2, 8  
 Belczynski, K., Sadowski, A., Rasio, F. A., & Bulik, T. 2006, *ApJ*, 650, 303  
 Binney, J., & Tremaine, S. 2011, *Galactic Dynamics* (2nd ed.; Princeton, NJ: Princeton Univ. Press)  
 Breen, P. G., & Hoggie, D. C. 2013, *MNRAS*, 432, 2779  
 Capuzzo-Dolcetta, R., Spera, M., & Punzo, D. 2013, *JCoPh*, 236, 580  
 Chomiuk, L., Strader, J., Maccarone, T. J., et al. 2013, *ApJ*, 777, 69  
 Clark, G. W. 1975, *ApJ*, 199, L143  
 Dehnen, W., & Read, J. I. 2011, *EPJP*, 126, 28  
 Fregeau, J. M., Cheung, P., Portegies Zwart, S. F., & Rasio, F. A. 2004, *MNRAS*, 352, 1  
 Fregeau, J. M., Joshi, K. J., Portegies Zwart, S. F., & Rasio, F. A. 2002, *ApJ*, 570, 171  
 Freitag, M. 2008, *Lecture Notes in Physics*, Vol. 760, The Cambridge N-Body Lectures, ed. S. J. Aarseth, C. A. Tout, & R. A. Mardling (Springer Netherlands), 123–158, doi:10.1007/978-1-4020-8431-7  
 Freitag, M., & Benz, W. 2001, *A&A*, 375, 711  
 Gaburov, E., Harfst, S., & Zwart, S. P. 2009, *New A*, 14, 630  
 Giersz, M. 1998, *MNRAS*, 298, 1239  
 Harfst, S., Gualandris, A., Merritt, D., & Mikkola, S. 2008, *MNRAS*, 389, 2  
 Hoggie, D. C. 2014, *MNRAS*, 445, 3435  
 Hénon, M. 1971, *Ap&SS*, 14, 151  
 Hurley, J. R., Pols, O. R., & Tout, C. A. 2000, *MNRAS*, 315, 543  
 Hurley, J. R., Tout, C. A., Aarseth, S. J., & Pols, O. R. 2001, *MNRAS*, 323, 630  
 Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, *MNRAS*, 329, 897  
 Joshi, K. J., Rasio, F. A., Zwart, S. P., & Portegies Zwart, S. 2000, *ApJ*, 540, 969  
 Maccarone, T. J., Kundu, A., Zepf, S. E., & Rhode, K. L. 2007, *Nature*, 445, 183  
 Mackey, A. D., Wilkinson, M. I., Davies, M. B., & Gilmore, G. F. 2008, *MNRAS*, 386, 65  
 McMillan, S. L. W. 1986, *ApJ*, 307, 126  
 McMillan, S. L. W., & Lightman, A. P. 1984a, *ApJ*, 283, 813  
 —. 1984b, *ApJ*, 283, 801  
 Miller-Jones, J. C. A., Strader, J., Heinke, C. O., et al. 2015, *MNRAS*, 453, 3919  
 Morscher, M., Pattabiraman, B., Rodriguez, C., Rasio, F. A., & Umbreit, S. 2015, *ApJ*, 800, 9  
 Morscher, M., Umbreit, S., Farr, W. M., & Rasio, F. A. 2013, *ApJ*, 763, L15  
 Nitadori, K., & Aarseth, S. J. 2012, *MNRAS*, 424, 545  
 Nitadori, K., & Makino, J. 2008, *New A*, 13, 498  
 Pattabiraman, B., Umbreit, S., Liao, W.-k., et al. 2013, *ApJS*, 204, 15  
 Pooley, D., Lewin, W. H. G., Anderson, S. F., et al. 2003, *ApJ*, 591, L131  
 Portegies Zwart, S. F., McMillan, S. L. W., Hut, P., & Makino, J. 2001, *MNRAS*, 321, 199  
 Rodriguez, C. L., Morscher, M., Pattabiraman, B., et al. 2015, *Phys. Rev. Lett.*, 115, 051101  
 Sigurdsson, S., & Hernquist, L. 1993, *Nature*, 364, 423  
 Spitzer, L. J. 1969, *ApJ*, 158, L139  
 Stodoikiewicz, J. S. 1982, *Acta Astron.*, vol. 32, no. 1-2, 1982, p. 63-91.  
 Strader, J., Chomiuk, L., Maccarone, T. J., Miller-Jones, J. C. A., & Seth, A. C. 2012, *Nature*, 490, 71  
 Strader, J., Seth, A. C., Forbes, D. A., et al. 2013, *ApJ*, 775, L6  
 Vasiliev, E. 2014, *MNRAS*, 446, 3150  
 Vasiliev, E., Antonini, F., & Merritt, D. 2015, *ApJ*, 810, 49

Wang, L., Spurzem, R., Aarseth, S., et al. 2015, MNRAS , 450,  
4070